The ALMA high speed optical communication link is here. An essential component for reliable present and future operations.

G.Filippi*[a,b], J.Ibsen [a,b], S.Jaque[c], F.Liello[d], N.Ovando [a], A.Astudillo[c], J.Parra [a],Christian Saldias [a]
[a]Joint ALMA Observatory[1], Alonso de Córdova 3107, Vitacura, Santiago de Chile, Chile;
[b]European Organisation for Astronomical Research in the Southern Hemisphere (ESO[2]), Karl-Schwarzschild-Strasse 2, D-85428 Garching, Germany;
[c]Corporación REUNA[3], Canadá 239, Providencia, Santiago de Chile, Chile;
[d]Consortium GARR[4], Via dei Tizii, 6, I-00185 Roma, Italy.

ABSTRACT

Announced in 2012, started in 2013 and completed in 2015, the ALMA high bandwidth communication system has become a key factor to achieve the operational and scientific goals of ALMA. This paper summarizes the technical, organizational, and operational goals of the ALMA Optical Link Project, focused in the creation and operation of an effective and sustainable communication infrastructure to connect the ALMA Operations Support Facility and Array Operations Site, both located in the Atacama Desert in the Northern region of Chile, with the point of presence of REUNA in Antofagasta, about 400km away, and from there to the Santiago Central Office in the Chilean capital through the optical infrastructure created by the EC-funded EVALSO project and now an integral part of the REUNA backbone. This new infrastructure completed in 2014 and now operated on behalf of ALMA by REUNA, the Chilean National Research and Education Network, uses state of the art technologies, like dark fiber from newly built cables and DWDM transmission, allowing extending the reach of high capacity communication to the remote region where the Observatory is located. The paper also reports on the results obtained during the first year and a half testing and operation period, where different operational set ups have been experienced for data transfer, remote collaboration, etc. Finally, the authors will present a forward look of the impact of it to both the future scientific development of the Chajnantor Plateau, where many installations area are (and will be) located, as well as the potential Chilean scientific backbone long term development.

Keywords: high speed network, EVALSO, DWDM, fiber, high performance network, high capacity network, communication system, WAN, MAN.

1. INTRODUCTION

Astronomy sites are often chosen to be far away from human settlements to reduce man-induced interferences, like light and radio contamination. Very often this also means that the pristine sites have limited or null communication infrastructures, though availability of networking is a key success factor for nowadays observatories.

This situation was true also for the ALMA Observatory, located in the Northern region of Chile, and the optical link project started in 2012 and completed in 2015 managed to close the gap providing a Gbps class link between the observation facilities and the central offices, in Santiago de Chile. The project builds on the existing EVALSO [7] infrastructure that provides similar fiber base connectivity to, among others, the ESO Paranal Observatory.

Section 2 provides some general information on the ALMA Observatory and terms that are then used later in the paper.

Section 3 provides an overall view on project goals, architecture of the communication system, project timeline, and the construction phase.

Section 4 provides statistics, monitoring strategies, and experiences from the first period of use, from early 2015 to date.

Section 5 provides some examples on how the presence of the high capacity link changed and it is expected to change further the processes at the ALMA Observatory.

Section 6 provides links to other projects and possible future developments.

*gfilippi@eso.org; phone +56 (0)9 6646 8368; www.eso.org, www.almaobservatory.org

The optical link project is part of the ALMA Development program that, after the completion of the construction phase of the ALMA Observatory, provides the framework to enhance and further develop the ALMA scientific and technical capabilities beyond the construction phase. National Radio Astronomy Observatory (NRAO, [6]), under AUI Associated Universities Inc. (AUI [5]) management, is providing the bulk funding for the procurement of the new items between Antofagasta and the ALMA Observatory site. ESO contributed in kind with the project management and, thanks to its participation in the EVALSO project, indirectly contributes via REUNA [3] the long haul Antofagasta-Santiago.

2. THE ALMA OBSERVATORY

The Atacama Large Millimeter/submillimeter Array (ALMA), an international partnership of Europe, North America and East Asia in cooperation with the Republic of Chile, is the largest astronomical project in existence. ALMA is a single telescope of revolutionary design, composed of 66 high precision antennas located on the Chajnantor plateau, 5000 meters altitude in northern Chile.

On behalf of the three Executives from East Asia (EA), Europe (EU), and North America (NA) the Joint ALMA Observatory (JAO) [1] provides the unified leadership and management of the construction, commissioning and operation of ALMA.

The ALMA Observatory counts on three physical locations. Two of them are in the area close to San Pedro de Atacama and constitute the observing site. They are (see Figure 1):

Figure 1 ALMA sites in Chile

- The ALMA Array Operations Site (AOS), at the Chajnantor plateau, 5000 meters, covering an area of approximately 15km radius;
- The ALMA Operations Support Facility (OSF), at 3000m, some 30 km southern of San Pedro de Atacama.

OSF and AOS are approximately 35 km apart and are connected by means of a private road, as well as by infrastructures for power and communication (fiber cable).

The third location **Error! Reference source not found.**is the ALMA Santiago Office (SCO), located in the same campus where ESO Offices are, in the Vitacura district of the Chilean capital. SCO has more than 100 offices and hosts the primary ALMA data archive.

The ALMA structure is completed by the ALMA Regional Centers (ARCs), one per Executive and located in the corresponding region in the Northern Hemisphere.

Data are produced at AOS-OSF, transferred to SCO and from there replicated to each of the ARCs. Data are stored in four ALMA Archives, one at SCO and one in each of the ARCs.

3. THE GENESIS AND CONSTRUCTION OF THE LINK

Following a feasibility study in 2011, the project proposal was endorsed by the ALMA governing body in April 2012. The procurement was completed in 2012. The construction ended on December 2014 with the first successful end-to-end connection between the ALMA AOS and the ALMA SCO.

3.1 Key goals

The key goals that were defined at the beginning were:

- Be a long term (>15 years) solution;
- Coping with projected operational needs (i.e., a Gbps-class infrastructure);
- Being available as soon as the array enters in full operation;
- Have Low Operational Costs (OPEX).

3.2 Project Architecture

To fulfill these, the analysis of the possible options for the area indicated the following solution (Figure 2):

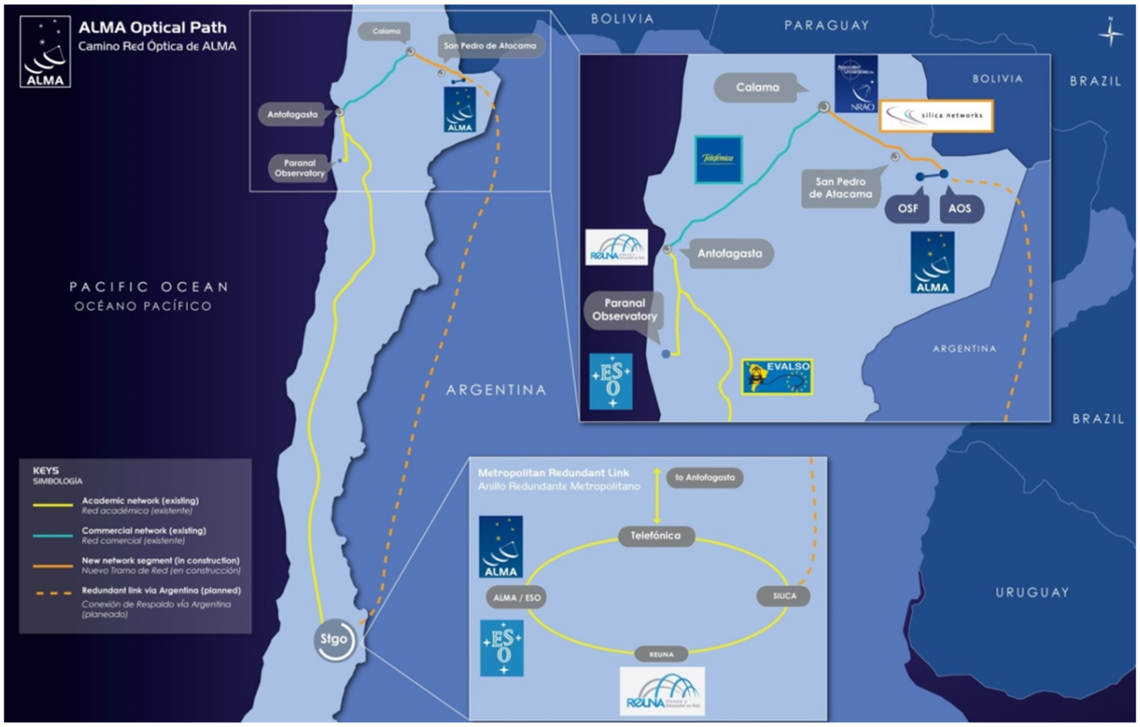

Figure 2: The overall topology of the (existing and) new communication system.

1. **A dark fiber pair between AOS and CALAMA** (about 150km) coming from a new cable installed by the company SILICA DATCO. This was completed in December 2014.
2. **A dedicated LAMBDA between CALAMA and the REUNA Point of Presence (PoP) in ANTOFAGASTA** (about 200 km) procured from Telefónica. This was delivered in December 2013
3. **A dedicated sub-LAMBDA between the REUNA PoP in ANTOFAGASTA and the SCO at the Vitacura Campus in Santiago** that is indeed configured on the existing EVALSO backbone that is supporting the Chilean academic network.
4. **DWDM equipment** (from PADTEC) **in ALMA/AOS, CALAMA, and ANTOFAGASTA.** The installation and integration with the physical links was done in the second half of 2014.
5. **The operation of the overall structure from the integrated NOC of REUNA.**

The result is an uninterrupted end-to-end 2.5 Gbps channel at exclusive use of the ALMA project.

3.3 Construction

Although a more detailed description of the early phases of the project and of the construction up to mid-2014 can be found in [13] and not repeated here, it is worth to report that in the second half of 2014 the construction of the cable faced the most challenging obstacles: the 3000m high mountain ridge between Calama and San Pedro de Atacama, the crossing of several roads and two river beds, some minefields (!), and last, but not least, the final climbing to the 5000 meters of the Chajnantor Plateau where the ALMA AOS is. In Figure 3 the layout of the new fiber cable that was laid down by the company SILICA DATCO. Contrary to the first part done in the first half of 2014, where a trenching machine could be used, the second part was built using mainly excavating machines. In Figure 4 some pictures of the construction of the second half of the fiber cable.

Figure 3: layout of the new fiber cable between Calama and ALMA AOS
(“Camino a”→ “Road to”; “Rio”→ “River”; “Campo Minado” → “Minefield”)

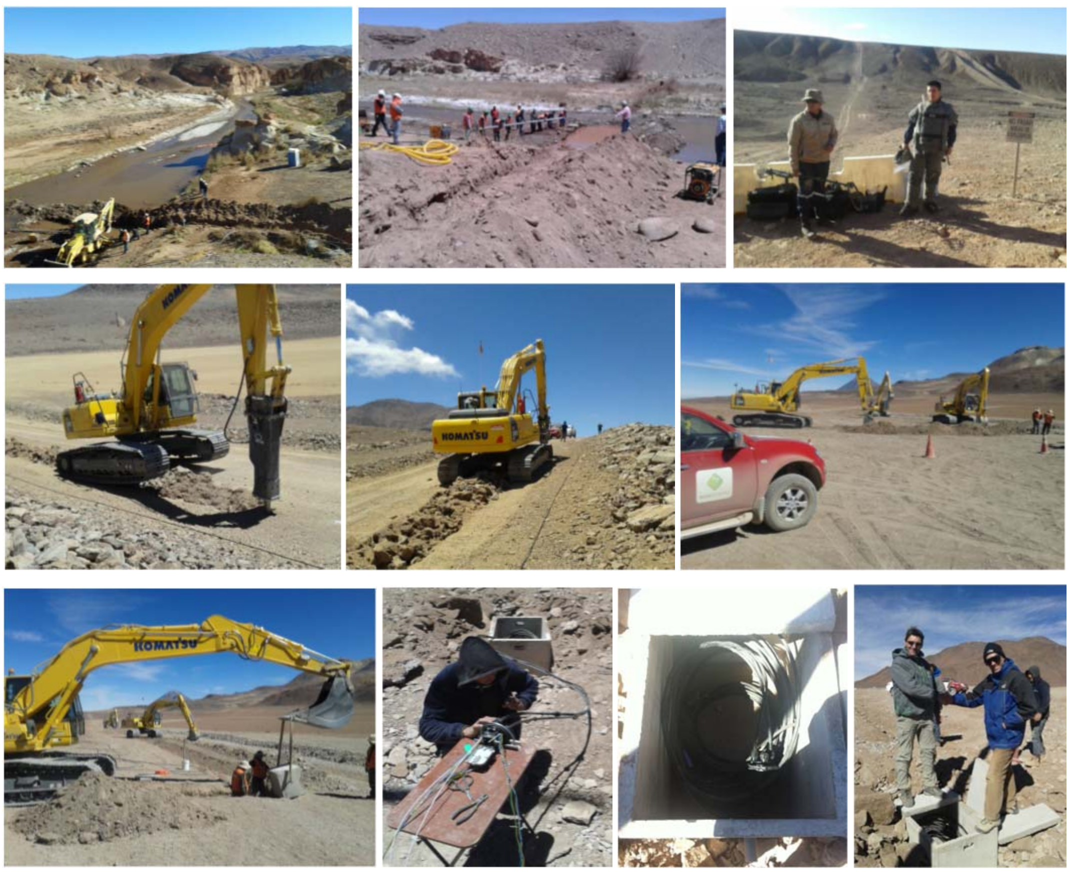

Figure 4: Various moments of the construction of the fiber cable. From the top left, crossing of the San Pedro and Vilama rivers; specialists from the Chilean Army clearing one of the minefields; opening the trench for the cable and installing the chamber for the splice box; final splicing of the ALMA fiber pair into the connection to the Observatory; final handshaking between the SILICA and ALMA Project Managers.

3.4 End-to-end testing

The end-to-end link was firstly tested on December 30, 2014, with file transfers and videoconference links between ALMA AOS and ALMA SCO (in Santiago). The results (see Figure 5) indicated that the designed capacity, namely two 1Gbps links out of the 2.5Gbps sub-lambda, could be fully exploited.

Report for **(Gi2/2/20) LINK-SCO-AOS**

Report Time Span:	12/30/2014 16:15:00 - 12/30/2014 16:45:00			
Sensor Type:	SNMP Traffic 64bit (30 s Interval)			
Probe, Group, Device:	Cluster Probe > Computer Room > N-DR-01-VSS.aos			
Cluster Node:	**SCO - PRTG**			
Uptime Stats:	Up:	100 % [29m30s]	Down:	0 % [0s]
Request Stats:	Good:	100 % [60]	Failed:	0 % [0]
Average (Traffic In):	883 Mbit/s			
Total (Traffic In):	194,095,756 KByte			

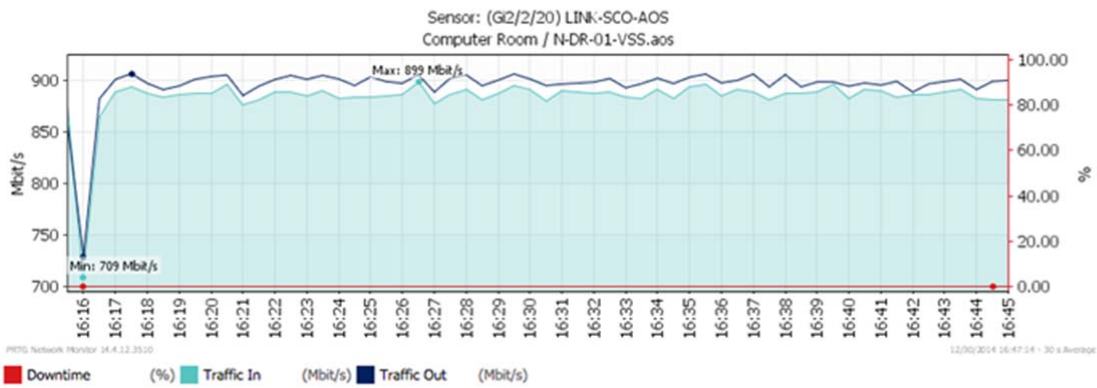

Figure 5: first tests of the end-to-end link on December 30, 2014.

After this initial testing, done on standalone architecture, the link has been integrated with the ALMA network architecture to become the main backbone for the communication between the Observatory sites, OSF and AOS, and the central office (SCO) in Santiago.

3.5 Has the Construction ended?

Unfortunately it is not yet possible to put the word end to the construction phase. Due to some delay in the handling of the official permits, at the time of writing, the system has not been formally delivered. It is common hope this may happen by the end of 2016. Anyhow thanks to the agreement between all parties, the system has been made available for extensive testing, and the result of this is reported in the following section.

4. OPERATING THE LINK (NETWORK VIEW)

The operation of the link is done by REUNA, the Chilean Research and Academic Network, from the Network Operation Center (NOC) in Santiago. In addition to this, ALMA Network group also keeps monitoring the overall communication system.

From the ALMA Network point of view, the available capacity of the link (2.5Gbps) has been configured as two independent 1Gbps links assigned as follows:

SCI Traffic: 1Gbps for scientific data only (Archive traffic)

ALMA Traffic: 1Gbps for the remaining non-scientific data traffic

The two new links and the previous ENTEL Microwave (0.1 Gbps) are bundled in a redundant configuration where in case of failure of one link the traffic is passed to another one, more specifically:

- if one of the two 1Gbps is not available, then the traffic is moved to the other 1Gbps one;
- if both are not available, then the traffic is moved to the Microwave link

This section provides elements on both the pure link (REUNA) as well as the ALMA networking point of view. As reality is also the unexpected, to complete the operational view, some dramatic events that affected the system are reported.

4.1 Monitoring the physical link (REUNA view)

REUNA NOC is in charge of monitoring the status of all the components of the link. This includes both the optical segments, dark fibers or lambdas, and the DWDM equipment operating the OTU2 levels and interfacing with the TELCO providers. The monitoring points are shown in (Figure 6).

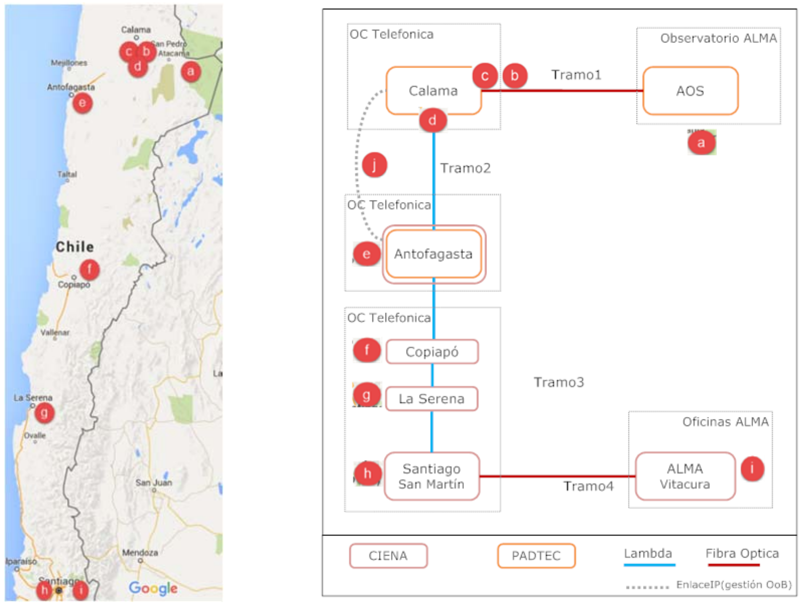

Figure 6 Monitoring points along the ALMA AOS-SCO link

To cover al lthe aspects, a mix of technologies and tools is in use, include the usage of Out of Band (OoB) tecniques.

The path AOS-Calama-Antofagasta is monitored using the Padtec monitoring platform, using a central node which collect the information, this node can be setup in a high availability schema, with a primary and secondary node. Together with this node there are the visualization nodes, these run over a java application.

The path Antofagasta-Santiago is monitored using the REUNA's NOC monitoring tool, based on a collection of open source and some comercial tools, the statistics are collected by SNMP polling the DWDM equipment in the path.

The main tools used are: Big Brother for reachability, MRTG for traffic statistics and PRTG for receiving optical Power.

Reachability

Big Brother (is configured to fire a ping test once every 3 minutes. If the Ping is missing the tool display an alarm. This usually means that there's a problem with the link or the remote equipment.

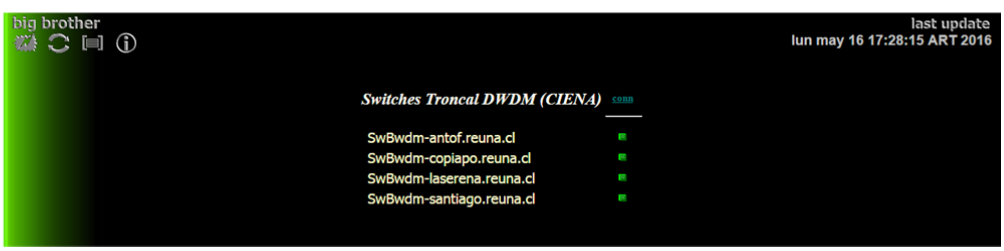

Figure 7 Big Brother used to monitor Reachability

Traffic Statistics

This information is gathered using SNMP and MRTG to graph. Figure 8 provide an example.

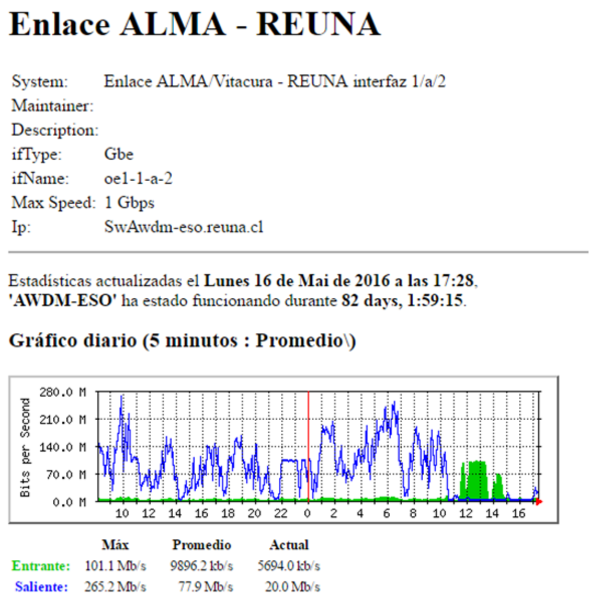

Figure 8 Example of PRTG graph

Optical Power

The optical power is the strength of the signal that an interface is receiving from the remote side. The data are stored and graphically presented to have an historical record about how the physical links behaves over time. This is particularly useful to check how well the fixes are done after a fiber cut. The data are obtained from the DWDM units along the path using SNMP and then displayed using PRTG. Figure 9 provide the view over the instant Optical Power and an example of a single value over time.

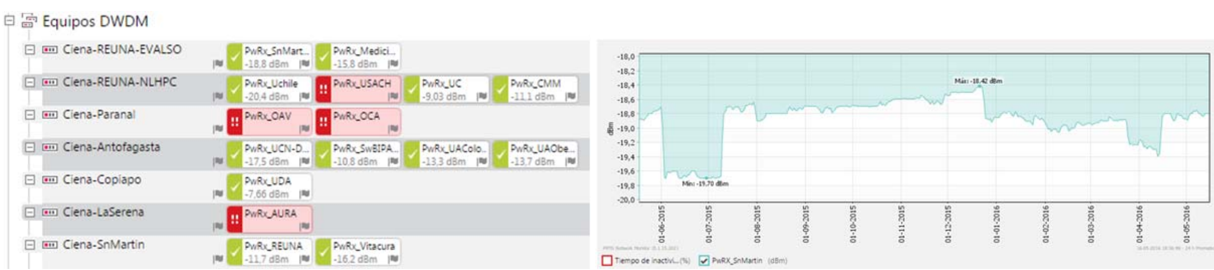

Figure 9 Nodes monitored for Optical Power and one example of display over time.

Functional Monitoring at IP level

In addition to the monitoring done on the optical links and on the DWDM equipment, in collaboration with ALMA Network team, REUNA also runs a functional monitoring based on the use of a VLAN connecting two points, one in the ALMA AOS and the second at REUNA. This allows having the health status not only of the AOS-SCO link, but also of the border routers and of the other layers needed for an application to application communication.

The display of the status is also done using the Big Brother tool, as shown in Figure 10.

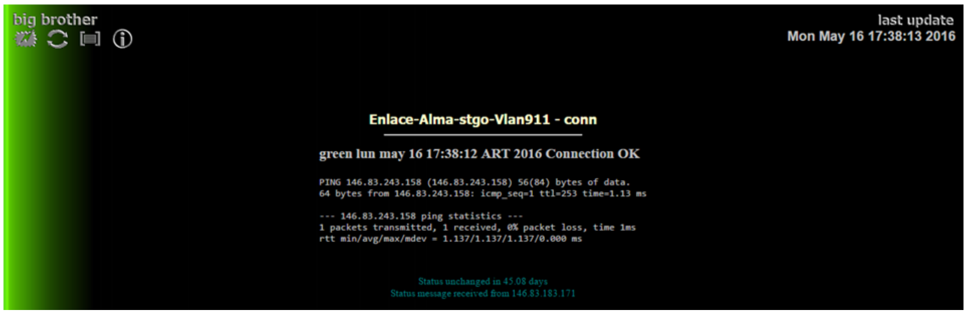

Figure 10 monitoring of the end-to-end functionality at IP level.

Reporting

Last but not least, REUNA is also in charge to provide monthly reports on the availability of the single parts and of the overall link, as well as the details of each failure. This provides the basis to handle the SLA with the different providers.

Examples taken from the April 2016 Report of availability over the month and accumulated are shown in Figure 11

Periods:

- Month: April 2016, 30 days, 720 hours
- Trimester: February 2016 to April 2016
- Year: September 2015 to April 2016, 243 days, 5832 hours

Service	Monthly			Yearly		
	Uptime [hours]	Downtime [hours]	Availability [%]	Uptime [hours]	Downtime [hours]	Availability [%]
1	720	2,16	97,30%	5832	53,01	99,09%

Table 3: Availability

3 Monthly Downtime

Item	Description	Month [hours]	Penalty Month [%]	Max. Unavailability [hours]
1	AOS – Calama fiber	0	0%	7,2
2	Lambda Calama - Antofagasta	0	0%	1
3	Lambda Antofagasta-Santiago	2,26	-	16
4	Santiago-ESO fiber	0	-	16
5	Santiago-REUNA fiber	0	-	16
6	DWDM equipment (Padtec)	0	0%	7,2

Table 4: Maximum unavailability

Figure 11 examples of statistics in the monthly report.

4.2 Monitoring of the data flow (ALMA view)

The ALMA network has been reconfigured to take full advantage of the new link. The first year of use of the link has confirmed the initial expectations in terms of bandwidth availability and reduced and stable latency:

- Delay (PING) between OSF and SCO (for both links is the same), around 24msec.
- Science Data Traffic: between 50 and 100 Mbps, and peaks between 200 and 400Mbps.
- Other ALMA Traffic: around 50 Mbps, with daily peaks to 100-130Mbps.

With the new architecture, the network team has more flexibility in the configuration in terms of multiple VLANs and VRFs, allowing more private networks. With higher available bandwitdh, the tranfer time of the scientific data shortened considerably, allowing easier and safer operations. Also quality of Videoconferencing and IP Telephony improved significantly.

To show the integration of the new and the old (now backup) links, Figure 12 shows the evolution of the average latency between AOS and SCO. Normally the value is around 24ms, that is the typical value for the new optical link, menaing the link was in use. In case of unavailability the traffic is automatically re-routed to the backup link, the 100Mbps microwave one, for which the delay is in the order of 50ms or higher. The downtime indicated situations where both primary and backup link were not available.

Report Time Span:	06/01/2015 00:00:00 - 05/01/2016 00:00:00				
Sensor Type:	Ping (30 s Interval)				
Probe, Group, Device:	Cluster Probe > Computer Room > N-DR-01-VSS.aos				
Cluster Node:	**SCO - PRTG**				
Uptime Stats:	Up:	99.997 % [334d14h42m23s]	Down:	0.003 % [12m6s]	
Request Stats:	Good:	99.912 % [963609]	Failed:	0.088 % [852]	
Average (Average):	26 msec				

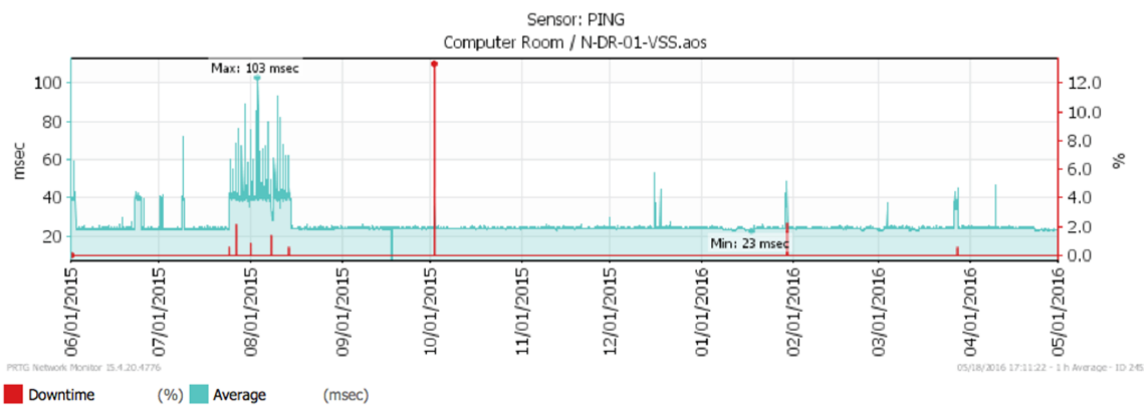

Figure 12 monitoring of the AOS-SCO latency in the period June2015-May2016

4.3 Relevant events

Beside the normal hiccups that equipment and procedures can experience during the initial period, there have been three major external events that affected the availability if the link:

- During the week March 23-27, 2015, the northern regions of Chile, and especially the area between Calama and La Serena, have experienced exceptional bad weather conditions with heavy rains that created floods and destroyed communication lines (road, fibers, etc.) and entire areas of towns with death toll and great damages to public and private infrastructures. This resulted on a complete cut of all type of communication between the northern and central part of Chile for several hours. The heavy rain also exposed the new fiber cable in several points (one example in Figure 13). The damages were later repaired.

Figure 13 fiber cable section exposed by the rain.

- Due to a vandalistic act during demonstrations in the center of Santiago, a key communication vault was set on fire on May 28. The fire caused the completed destruction of the cables present in the vault (Figure 14) including those part of the EVALSO infrastructure. Due to the amount of cables to be repaired by replacing the complete damaged sections, the repair work lasted several days and the service was back only on June 02.

Figure 14 Burned cable in a vault in Santiago city center.

- In July, a false maneuvering of a heavy truck damaged the concrete lid of a chamber in Calama (see Figure 15), causing damages to the fiber cores inside the cable. After a temporary fix, a new chamber was built some meters away, to reduce the risk of similar events in the future.

Figure 15 damaged chamber and fiber cable in Calama

5. USING THE LINK (APPLICATION VIEW)

The availability of the new link it is not only provides a better communication layer by and large, but it also allows reconsidering some application deployment that could be reorganized to better take advantage of the new performance in terms of bandwidth and latency. In addition to the normal IT services, like Videoconferencing, IP telephony, e-mail, etc., there are two major areas that have been already investigated: the Scientific Archive and the Remote Operations.

5.1 Scientific Archive

Because of the communication constraints of the previous communication architecture based on a 100Mbps link between AOS/OSF and SCO, the design of the Scientific Archive was based on two nodes, one at OSF and one at SCO that needed bidirectional replication. This is the current architecture in use.

In this configuration, a first limit is the capability of the link to transfer the data from the node (at OSF) that receives the data from the correlator to the node (at SCO) that acts as the main storage for later distribute the data to the regional centers (ARCs). Even though the 100Mbps is just enough to transfer ALMA data from OSF to SCO now (i.e., 100Mbps can transfer 1TB in 20 hours) without accumulating backlog, as ALMA operations is ramping up, soon this will not be the case. It is expected a maximum of 850GB/day in cycle4 and ˜1TB/day in cycle5, while in the current cycle3 the level is around 300-400 GB/day. This alone would be a justification to have the optical link, which will play a key role in data transmission from October 2017 onwards.

In addition to the obvious increase of the data transfer capability, the availability of the new link allows also the Data Archive team to consider replacing the current dual node scheme with a single database in SCO, with the obvious savings in terms of hardware and operational costs, especially when one is thinking that the data volume will increase in the future. At AOS/OSF there will be only a buffering stage, in case of unavailability of the link. In case of link cut, the high capacity of the new link would allow to recover rapidly the situation as soon as the cut is fixed.

The current and proposed layout is in Figure 16.

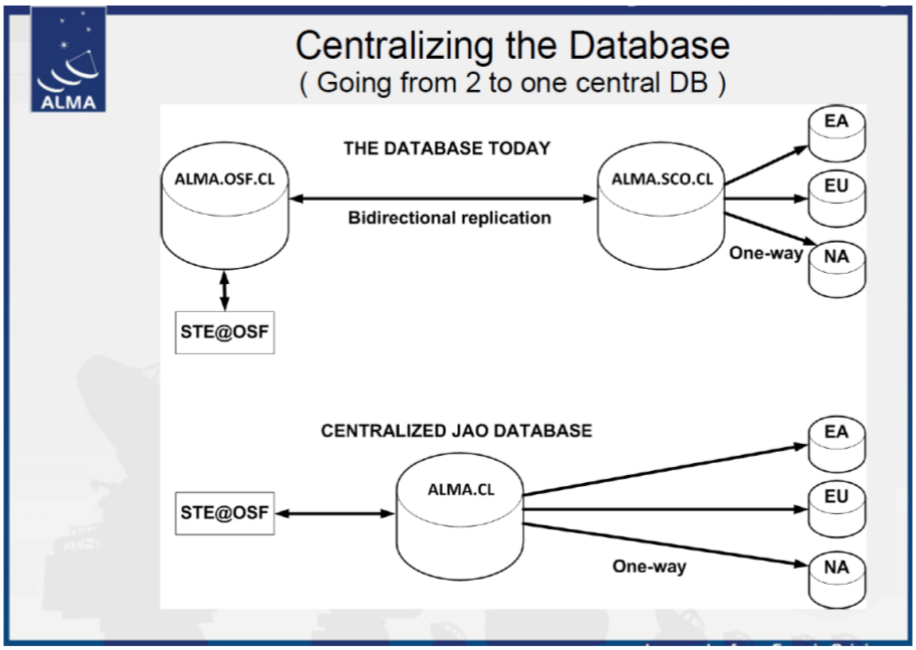

Figure 16 Current and proposed architecture for the Scientific Archive (Chile)

In order to support the new design, the Data Archive team has done some application testing to compare the performance of the new link against the old one (100Mbps) and using the performance of a LAN as baseline. In Figure 17 the result of operating some commonly used functions in the three scenarios. The performance over the new link (between SCO and OSF) is considered acceptable for all cases in comparison to the performance when the two peers are on the same LAN.

Process	LAN	1Gb	100Mb	Acceptable?
TMCDB loading	3 min.	13 min.	34 min.	✗
Create Array	50 sec.	2.5 min	5 min.	✗
Select Project	5 sec.	10 sec.	20 sec.	✓
Select Scheduling Block (SB)	1 sec.	1 sec.	1 sec.	✓
Confirm array for selected SB	1 sec.	2 sec.	2 sec.	✓
NGAS process/medium file	384 msec.	467 msec.	700 msec.	✓
Destroy Array	1 min.	1 min.	1 min.	✓

Figure 17 response time of some commonly used functions

Also the latency, that for this type of applications plays a pivotal role, is compatible with the application requirements, as shown in Figure 18.

- Feasibility – Online/Offline Systems separated from DB by:

Network	Speed	Latency	Online / Offline
LAN	1 Gb	0.5 ms	OK / OK
Fiber Link	1 Gb	23 ms	OK */ OK (3 times slower)
MW Link	100 Mb	30 ms	NOT OK / OK (10 times slower)

Figure 18 Latency measured between two nodes in different network configurations.

5.2 Remote Operations

Another area where the availability of high bandwidth and low and constant latency could be an asset is the remote operations of the Observatory. Already now the control room is at 30km (at OSF) from the antennas installation (AOS). A good link would allow bringing this activity to SCO, 1200km away, still with comparable performances, but lower costs and better comfort for the operators.

This option was originally explored and reported in [12] and will be further investigated in the coming years.

6. A LOOK INTO THE FUTURE

As soon as the administrative impasse is over, the ALMA optical link will enter its full operational life, and this will close the page. But there is still a book to be written.

REUNA in the frame of its infrastructure strategy plan is encouraging to extend the national backbone along the Chilean territory based on direct access to fiber, as the Calama – AOS path, and to move to a concept of “unlimited” bandwidth capacity backbone, as well as to keep the international links thanks to the collaboration with The Latin America Academic Network, CLARA (Cooperación Latinoamericana de Redes Avanzadas) [8], better known by RedCLARA; with its peering to, among others, the European GEANT and the North American INTERNET2, and the AmLight initiative [9].

In the national development plan, thanks to the agreement with the LSST project and AURA, during 2016 a path of 700Kms of dark fiber will be deployed (lighted with 100Gbps lambdas) to increase the capability of the REUNA backbone to support the research and academic community.

Also, with the goal to extend the dark fiber backbone to Calama and beyond, REUNA is one of the members of the BELLA initiative, which seek to extend a terrestrial network based on fiber along Latin America, fiber that will be connected to the new submarine cable from Europe to Latin America EULALINK [10].

7. CONCLUSIONS

High-bandwidth communication is a key factor for scientific installations as Observatories and the ALMA optical link is delivering the promised value. Beyond purely increasing the bandwidth and lowering latency, boosting performances of applications directly making use of these two parameters, the new system is also an enabler for potential processes transformations that could deliver further advantages.

The project, now about to enter the full exploitation, it is also part of a more general evolution of the academic and research network in Chile that in the coming years could bring even more advantages to ALMA and to the other scientific projects in the area.

REFERENCES

[1] http://www.almaobservatory.org/

[2] http://www.eso.org/

[3] http://www.reuna.cl/

[4] http://www.garr.it/

[5] /http://www.aui.cl

[6] http://www.nrao.edu

[7] http://www.evalso.eu/

[8] http://www.redclara.net

[9] http://www.amlight.net/index.html

[10] http://www.ella-int.eu/index.php/2-home-page/1-the-ella-project

[11] Filippi, G. et Al, "EVALSO: a high-bandwidth communication infrastructure to efficiently connect the ESO Paranal and the Cerro Armazones Observatories to Europe" SPIE 2010, SPIE 7740-53, (2010)

[12] Shen T., Soto R., Ibsen J. et Al, "Exploring Remote Operation for ALMA Observatory" SPIE 2014, SPIE 9149-90, (2014)

[13] Filippi, G. et Al, " ALMA communication backbone in Chile goes optical” SPIE 2014, SPIE 9152-29, (2014)